

FORM PTO-1390
(REV. 1-98)

U.S. DEPARTMENT OF COMMERCE PATENT AND TRADEMARK OFFICE

ATTORNEY'S DOCKET NUMBER

TRANSMITTAL LETTER TO THE UNITED STATES
DESIGNATED/ELECTED OFFICE (DO/EO/US)
CONCERNING A FILING UNDER 35 U.S.C. 371

3696-0101P

U.S. APPLICATION NO. (If known, see 37 CFR 1.5)

09/743268
NEW

INTERNATIONAL APPLICATION NO.

PCT/NO00/00155

INTERNATIONAL FILING DATE

May 10, 2000

PRIORITY DATE CLAIMED

May 10, 1999

TITLE OF INVENTION

A SEARCH ENGINE WITH TWO-DIMENSIONAL LINEARLY SCALABLE PARALLEL ARCHITECTURE

APPLICANT(S) FOR DO/EO/US

SVINGEN, Børge; RISVIK, Knut Magne; HALAAS, Arne and EGGE, Tor

Applicant herewith submits to the United States Designated/Elected Office (DO/EO/US) the following items and other information:

1. ☒ This is a **FIRST** submission of items concerning a filing under 35 U.S.C. 371.
2. ☐ This is a **SECOND** or **SUBSEQUENT** submission of items concerning a filing under 35 U.S.C. 371.
3. ☒ This express request to begin national examination procedures (35 U.S.C. 371(f)) at any time rather than delay examination until the expiration of the applicable time limit set in 35 U.S.C. 371(b) and PCT Articles 22 and 39 (1).
4. ☐ A proper Demand for International Preliminary Examination was made by the 19th month from the earliest claimed priority date
5. ☒ A copy of the International Application as filed (35 U.S.C. 371(c)(2))
 - a. ☐ is transmitted herewith (required only if not transmitted by the International Bureau).
 - b. ☒ has been transmitted by the International Bureau. WO 00/68834
 - c. ☐ is not required, as the application was filed in the United States Receiving Office (RO/US).
6. ☒ A translation of the International Application into English (35 U.S.C. 371(c)(3)).
7. ☒ Amendments to the claims of the International Application under PCT Article 19 (35 U.S.C. 371(c)(2)).
 - a. ☐ are transmitted herewith (required only if not transmitted by the International Bureau).
 - b. ☐ have been transmitted by the International Bureau.
 - c. ☐ have not been made; however, the time limit for making such amendments has NOT expired.
 - d. ☒ have not been made and will not be made.
8. ☐ A translation of the amendments to the claims under PCT Article 19 (35 U.S.C. 371(c)(3)).
9. ☐ An oath or declaration of the inventor(s) (35 U.S.C. 371(c)(4)).
10. ☐ A translation of the annexes to the International Preliminary Examination Report under PCT Article 36 (35 U.S.C. 371(c)(5)).

Items 11. to 16. below concern document(s) or information included:

11. ☒ An Information Disclosure Statement under 37 CFR 1.97 and 1.98.-1449 and International Search Report w/ cited documents
12. ☐ An assignment document for recording. A separate cover sheet in compliance with 37 CFR 3.28 and 3.31 is included.
13. ☒ A **FIRST** preliminary amendment.
☐ A **SECOND** or **SUBSEQUENT** preliminary amendment.
14. ☐ A substitute specification.
15. ☐ A change of power of attorney and/or address letter.
16. ☒ Other items or information:
PCT/ISA/210
PCT/RO/101
Six (6) sheets of formal drawings

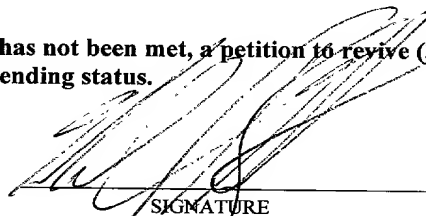
U.S. APPLICATION NO (if known, see 37 CFR 1.5) 09/743268		INTERNATIONAL APPLICATION NO PCT/NO00/00155		ATTORNEY'S DOCKET NUMBER 3696-0101P	
17. <input checked="" type="checkbox"/> The following fees are submitted:				CALCULATIONS PTO USE ONLY	
BASIC NATIONAL FEE (37 CFR 1.492(a)(1)-(5): Neither international preliminary examination fee (37 CFR 1.482) nor international search fee (37 CFR 1.445(a)(2)) paid to USPTO and International Search Report not prepared by the EPO or JPO.				\$1,000.00	
International preliminary examination fee (37 CFR 1.482) not paid to USPTO but International Search Report prepared by the EPO or JPO				\$860.00	
International preliminary examination fee (37 CFR 1.482) not paid to USPTO but international search fee (37 CFR 1.445(a)(2)) paid to USPTO.				\$710.00	
International preliminary examination fee (37 CFR 1.482) paid to USPTO but all claims did not satisfy provisions of PCT Article 33(1)-(4)				\$690.00	
International preliminary examination fee (37 CFR 1.482) paid to USPTO and all claims satisfied provisions of PCT Article 33(1)-(4)				\$100.00	
ENTER APPROPRIATE BASIC FEE AMOUNT =				\$ 1000.00	
Surcharge of \$130.00 for furnishing the oath or declaration later than <input checked="" type="checkbox"/> 20 <input type="checkbox"/> 30 months from the earliest claimed priority date (37 CFR 1.492(e)).				\$ 130.00	
CLAIMS	NUMBER FILED	NUMBER EXTRA	RATE		
Total Claims	11 - 20 =	0	X \$18.00	\$	0
Independent Claims	1 - 3 =	0	X \$80.00	\$	0
MULTIPLE DEPENDENT CLAIM(S) (if applicable) none			+ \$270.00	\$	0
TOTAL OF ABOVE CALCULATIONS =				\$ 1130.00	
Reduction of 1/2 for filing by small entity, if applicable. Applicant claims Small Entity Status in accordance with 37 CFR 1.27.				\$ 0	
SUBTOTAL =				\$ 1130.00	
Processing fee of \$130.00 for furnishing the English translation later than <input type="checkbox"/> 20 <input type="checkbox"/> 30 months from the earliest claimed priority date (37 CFR 1.492(f)).				\$ 0	
TOTAL NATIONAL FEE =				\$ 1130.00	
Fee for recording the enclosed assignment (37 CFR 1.21(h)). The assignment must be accompanied by an appropriate cover sheet (37 CFR 3.28, 3.31). \$40.00 per property +				\$ 0	
TOTAL FEES ENCLOSED =				\$ 1130.00	
				Amount to be: refunded \$	
				charged \$	

- a. ☒ A check in the amount of \$ **1130.00** to cover the above fees is enclosed.
- b. ☐ Please charge my Deposit Account. No. _____ in the amount of \$ _____ to cover the above fees.
A duplicate copy of this sheet is enclosed.
- c. ☒ The Commissioner is hereby authorized to charge any additional fees which may be required, or credit any
overpayment to Deposit Account No. 02-2448.

NOTE: Where an appropriate time limit under 37 CFR 1.494 or 1.495 has not been met, a petition to revive (37 CFR 1.137(a) or (b)) must be filed and granted to restore the application to pending status.

Send all correspondence to:

Birch, Stewart, Kolasch & Birch, LLP or Customer No. 2292
P.O. Box 747
Falls Church, VA 22040-0747
(703)205-8000


SIGNATURE

MUTTER, MICHAEL K.
NAME

#29,680 (MKM)
REGISTRATION NO.

09/743268

534 Rec'd PCT/PTO 08 JAN 2001
PATENT
3696-0101P

IN THE U.S. PATENT AND TRADEMARK OFFICE

Applicant: SVINGEN, Børge et al.
Int'l. Appl. No.: PCT/NO00/00155
Appl. No.: NEW Group:
Filed: January 8, 2001 Examiner:
For: A SEARCH ENGINE WITH TWO-
DIMENSIONAL LINEARLY SCALABLE
PARALLEL ARCHITECTURE

PRELIMINARY AMENDMENT

BOX PATENT APPLICATION

Assistant Commissioner for Patents
Washington, DC 20231

January 8, 2001

Sir:

The following Preliminary Amendments and Remarks are respectfully submitted in connection with the above-identified application.

AMENDMENTS

IN THE SPECIFICATION:

Please amend the specification as follows:

Before line 1, insert --This application is the national phase under 35 U.S.C. § 371 of PCT International Application No. PCT/NO00/00155 which has an International filing date of May 10, 2000, which designated the United States of America.--

REMARKS

The specification has been amended to provide a cross-reference to the previously filed International Application.

If necessary, the Commissioner is hereby authorized in this, concurrent, and future replies, to charge payment or credit any overpayment to Deposit Account No. 02-2448 for any additional fees required under 37 C.F.R. § 1.16 or under 37 C.F.R. § 1.17; particularly, extension of time fees.

Respectfully submitted,

BIRCH, STEWART, KOLASCH & BIRCH, LLP

By 

Michael K. Mutter, #29,680

MKM/rem
3696-0101P

P.O. Box 747
Falls Church, VA 22040-0747
(703) 205-8000

(Rev. 04/19/2000)

534 Rec'd PCT/PTO 08 JAN2001

A search engine with two-dimensional linearly scalable parallel architecture.

The present invention concerns a search engine with two-dimensional linearly scalable parallel architecture for searching a collection of text documents D . wherein the documents can be divided into a number of partitions d_1, d_2, \dots, d_n , wherein the collection of documents D is preprocessed in a text filtration system such that a preprocessed document collection D_p is obtained and corresponding preprocessed partitions $d_{p1}, d_{p2}, \dots, d_{pn}$, wherein an index I can be generated from the document collection D such that for each previous preprocessed partition $d_{p1}, d_{p2}, \dots, d_{pn}$ a corresponding index i_1, i_2, \dots, i_n is obtained, wherein searching a partition d of the document collection D takes place with a partition-dependent data set $d_{p,k}$ comprising both the preprocessed partition d_{pk} and the corresponding index i_k , with $1 \leq k \leq n$, and wherein the search engine comprises data processing units which form sets of nodes connected in a network.

Most prior art search engines work with large data set and employ powerful computers to perform the search. However, searching is a partitionable data processing problem, and this fact can be used to partition a search problem into a large number of specific queries and let each query be processed simultaneously on a commensurate number of processors connected in parallel in a network. Particularly searching can be regarded as a binary partitionable data processing problem, and hence a binary tree network is used for establishing a multiprocessor architecture such as disclosed for instance in US patent No. 4860201 (Stolfo & al.) and international patent application PCT/NO99/00308 which belongs to the applicant and hereby is incorporated by reference. The present applicant has developed proprietary technologies for searching within regular text documents. These technologies are i.a. based on a search system and a method for searching as described in international patent application PCT/NO99/00233 which belongs to the applicant and hereby is incorporated by reference. The search system is based on efficient core search algorithms which may be used in the search engine according to the invention.

However, it has become increasingly important to cater for a growing number of documents to be searched and also to be able to handle an increased traffic load, i.e. the number of queries per second which shall be processed by the

search system. This, apart from the ability to handle a large number of queries simultaneously on processor level, implies that a search engine should be implemented with an architecture that allows for preferably linear scalability in two dimensions, viz. both with regard to the data volume and the performance, i.e. the ability to handle a very large number of queries per second. Considering the development of the World Wide Web, a scalability problem in the search engine architecture will be extremely important as there presently is an enormous growth rate in both the number of documents and the number of users on the Internet.

Prior art search engine solutions for Internet are able to scale to a certain level, but almost always this is achieved in a manner that requires a high cost increase of the search engine system relative to the growth in data volume or data traffic. Very often the system costs scale as the square of the data volume or the traffic, a doubling of the data volume thus leading to quadrupled system costs. Furthermore all the major Internet search engines presently are based on very expensive server technology, often coupled with brute computing force-approaches and accompanied with disadvantages such as slow server turnaround, requirements for special hardware to provide fault tolerance etc. The system costs can e.g. be measured as the amount of hardware required to implement a search engine solution or the actual aggregated price of the system.

It is thus a main object of the present invention to provide a search engine with a multilevel data and functional parallelism, such that large volumes of data can be searched efficiently and very fast by a large number of users simultaneously.

Particularly it is a further object of the invention to provide a parallel architecture for implementing a search engine with a multilevel data and functional parallelism.

Yet a further object of the present invention is to provide a parallel architecture which is linearly scalable in two dimensions, i.e. with regard to both data volume and performance, that is the query rate.

The above-mentioned objects and further features and advantages are provided with a search engine according to the invention which is characterized in that the first set of nodes comprises α dispatch nodes, a

second set of nodes comprises b search nodes, a third set of nodes comprises g indexing nodes, and an optional fourth set of nodes comprises e acquisition nodes, that the dispatch nodes are connected in a multilevel configuration in the network, that the search nodes are grouped in columns which are
5 connected in parallel in the network between the dispatch nodes and an indexing node, that the dispatch nodes are adapted to process search queries and search answers, the search queries being dispatched further to all search nodes and in case the acquisition nodes are not present, the search answers being returned to the dispatch nodes and therein being combined to a final
10 search result, that the search nodes each are adapted to contain search software, that at least some of the search nodes additionally comprise at least one search processor module, that the indexing nodes are adapted for generally generating indexes i for the search software and optionally for generating partition-dependent data sets $d_{p,k}$ to search nodes which comprise
15 a search processor module, that in case acquisition nodes are present, these are connected in a multilevel configuration in the network similar to that of the dispatch node, and adapted for gathering answers to search queries and outputting a final result thereof, thus relieving the dispatch nodes of this task, and that the two-dimensional linear scaling respectively takes place by
20 scaling of the data volume through an increase in the number of partitions d and scaling of performance through replication of one or more partitions d .

According to the invention are advantageously the multilevel configuration of the dispatch nodes and the optional acquisition nodes network provided by hierarchical tree structures, and the multilevel configuration of the optional
25 acquisition nodes is then preferably a mirror image of the multilevel configuration of the dispatch nodes, the hierarchical tree structures preferably being binary tree structures.

According to the invention each of the search nodes advantageously comprises a search software module.

30 Further, according to the invention at least some of the search nodes comprises at least one dedicated search processor module, each dedicated search processor module being realized with one or more dedicated processor chips, which is adapted for parallel handling of a number of search queries. In this connection it is preferred that the dedicated search processor chips are
35 provided in the search processor modules in y processor groups, each with z

search processor chips and being connected with and adapted to receive data from a memory assigned to the processor group.

According to the invention the increase in the number of partitions in the scaling of the data volume is advantageously implemented by an increase in the number of search node groups or columns. In this connection the increase in the number of partitions can preferably be accompanied by a corresponding increase in the number of dispatch nodes and, in case, also in the number of acquisition nodes, and optionally also by an increase in the number of index nodes.

According to the invention the replication of one or more partition in the scaling of performance is advantageously implemented by an increase in the number of search nodes in each group or column.

Finally it is according to the invention advantageous that the separate node sets each is implemented over one or more workstations connected in a data communications network.

The search engine according to the invention shall now be described in terms of non-limiting exemplary embodiments and with reference to the accompanying drawings, in which

fig. 1 shows a general overview of the architecture in a first embodiment of the search engine according to the invention.

fig. 2 a general overview of the architecture in a second embodiment of the search engine according to the invention,

fig. 3 schematically a search node with search software and dedicated hardware search chips,

fig. 4 the layout of a single module of search chips.

fig. 5 schematically the query handling in a dedicated search chip,

fig. 6 schematically the principle of performance scaling,

fig. 7 a multilevel dispatch node configuration based on a binary tree structure.

fig. 8 an overview of the architecture of a first embodiment of the search engine according to the invention, with an indication of the principle for two-dimensional scaling, and

5 fig. 9 an overview of the architecture of a second embodiment of the search engine according to the invention, with an indication of the principle for two-dimensional scaling.

Searching a large collection of independent documents is a highly parallel task. The search engine according to the invention employs parallelism on different levels as shall be discussed in the following.

10 The search engine according to the present invention searches a document collection of documents D . The documents can be divided into n partitions d_1, d_2, \dots, d_n . Each document collection D , or partition d of a document collection can be preprocessed for use in a hardware text filtering system, for instance implemented by dedicated hardware like the applicant's so-called
15 Pattern Matching Chip (PMC) which is disclosed in the applicant's international patent application No. PCT/NO99/00344 which hereby is incorporated by reference. The preprocessed document collection is denoted D_p and the corresponding preprocessed document collection partitions d_p are denoted $d_{p1}, d_{p2}, \dots, d_{pn}$.

20 Software-based search systems require an index generated from the document collection. The index is denoted I and the indexes corresponding to the document collection partitions $d_{p1}, d_{p2}, \dots, d_{pn}$, are denoted i_1, i_2, \dots, i_n .

The data set needed for searching a partition d of the document collection D , is called the partition-required or partition-dependent data set. In a software
25 only system (SW system) the data set is the index i_k , while in the systems with hardware (SW/HW systems), the data set also includes the preprocessed document collection partition $d_{p,k}$ with the corresponding index i_k , where $1 \leq k \leq n$.

The essentially software-implemented partitioning and preprocessing
30 operations can be rendered schematically as

$*(D) \rightarrow *(d_1, \dots, d_n) \rightarrow *(d_{p1}, \dots, d_{pn}) \rightarrow d_{p,k}$, where $*(D)$ denotes a partitioning operation on the input D , $*(d_1, \dots, d_n)$ a filtering operation, e.g. indexing, on d_1, \dots, d_n and $d_{p,k}$ of course is the partition-dependent data set, which in an SW system only shall be the index i_k , and with $1 \leq k \leq n$.

A search engine is implemented on a cluster of workstations that are connected using a high performance interconnect bus. The not shown workstations then constitute the server of the search system (search server). The workstations implement nodes of the search server. The nodes perform different tasks and are according to the invention implemented as set out below.

The nodes can be regarded as virtual nodes distributed among the workstations. but in a SW/HW search engine the dedicated search processing hardware must be physically present in some workstations in order to support the hardware-based search nodes. The search node software may then still be distributed. Also, some search nodes in SW/HW search engine may comprise software only, and optionally be distributed over more than one workstation.

- Dispatch nodes $N_{\alpha 1}, \dots, N_{\alpha a}$ handle incoming queries and send queries along to all the search nodes. The dispatch nodes can also be configured as acquisition nodes, gathering the answers, i.e. the search results to the queries. Upon receiving the answers, the dispatch nodes in the acquisition mode merge the search results into a final result.
- Search nodes $N_{\beta 1}, \dots, N_{\beta b}$ hold a portion of the entire data set d_p . A search node comprises both the dedicated search software as well as a number of the above-mentioned PMC modules for hardware searching.
- Indexing nodes $N_{\gamma 1}, \dots, N_{\gamma g}$ are responsible for generating indexes for the dedicated search software on a number of search nodes. For the PMC modules, the indexing node N_{γ} also filters and preprocesses the raw data.
- Acquisition nodes $N_{\delta 1}, \dots, N_{\delta e}$ may optionally be provided for gathering the answers and merging the search results into a final result, in which case the dispatch node N_{α} , of course, is relieved of the acquisition task.

A first embodiment of the search engine according to the invention is shown in fig. 1, where the search nodes N_{β} are provided in columns or groups S provided and connected in parallel between the dispatch nodes N_{α} and the indexing nodes N_{γ} . The arrangement of the dispatch nodes N_{α} is shown schematically, but in practice they would be provided in a multilevel hierarchical arrangement.

A second embodiment of the search engine according to the invention is shown in fig. 2, where acquisition nodes N_s are provided in a similar arrangement as that of the dispatch nodes, relieving the latter of the acquisition task.

5 It is to be understood that singular workstations may implement a specific type of nodes only, or alternatively more than one type of nodes. In other words, the different types of nodes may be distributed over the cluster of workstations. Hence the architecture shown in figs. 1 and 2 is implemented by the whole cluster, and these figures accordingly do neither show the
10 workstations nor the interconnect bus.

The nodes shall now be discussed in more detail, starting with the search nodes which are central to the search engine according to the invention.

A search node N_β holds as mentioned a portion of the entire data set d_p . The search node has both a software search engine SW, and optionally a number
15 of PMC modules M, as shown in fig. 3. The data set for a search node N_β is generated on an indexing node N_γ , of which more anon.

A search node may be equipped with a number x of PMC modules M for very fast searching, as shown in fig. 3. Each PMC module M has y groups G of z PMCs, as shown in figure 4, where each group G receives data from a single
20 memory chip RAM. These A module M will typically be an individual circuit board. Each chip PMC is capable of processing q simultaneous queries, as shown schematically in fig. 5.

A pattern matching chip PMC can process a data volume of t_c bytes per second. Assuming that the memory modules are capable of delivering t_y bytes
25 per second to the pattern matching chips PMCs, a PMC can search through the data volume of T_c bytes, $T_c = \min\{t_c, t_y\}t$, in the given time t .

As shown in figure 4, the pattern matching chips PMC are placed in modules M with y groups G of z chips PMC, where each group G receives data from a single memory chip RAM, and the size of the memory chip is T_c . The total
30 amount of data this module can search through is $T_r = T_c$ with zq different queries.

When x modules M are provided in a search node N_β , these PMC modules M can search through an amount of data equal to $T_r = T_c \cdot x = \min\{t_c, t_y\}txy -$

since no PMC modules search through the same data, the number of concurrent queries is still zq .

Thus the total query rate of the PMC modules in a search node can be expressed as

$$5 \quad r_{HW} = \frac{zq}{\frac{T_r}{\min\{t_c, t_y\} \alpha \gamma}} \quad (1)$$

where T_r denotes the total data volume on a node. The search node performance can now be calculated.

Given that the PMC modules M (or any hardware equivalent) has a query rate of r_{HW} and that the search software on a search node N_β has a query rate of r_{SW} , the total query rate r_Σ of a search node N_s can be expressed as

$$10 \quad r_\Sigma = r_{t,HW} (1 - \varphi_{SW}) + r_{t,SW} \varphi_{SW} \quad (2)$$

where φ_{SW} denotes the percentage of queries q that will be executed in software. The actual value of φ_{SW} is dynamically updated at runtime from a statistical model.

15 The dispatch nodes N_α receive all the queries, and resend them to all the search nodes N_β . The answers from the different search nodes N_β are merged and in case the dispatch nodes N_α functions as acquisition nodes, a complete answer is returned.

20 The indexing nodes N_γ collect documents and create prebuilt indexes for the search software on the different search nodes N_β . Hence the indexing nodes N_γ can be incorporated in the search nodes N_β with appropriate indexing software in the latter. The hardware is based on scanning through the entire collection of raw data, but some preprocessing and filtering of the raw data can be done in the indexing nodes N_γ .

25 Concerning the interconnect and data traffic, some general observations can be made based on the following considerations.

Different types of interconnect can be used for connecting the nodes. For a lower end system, a regular 100 Mbit Fast Ethernet will for instance handle the traffic.

Traffic on the interconnect between the nodes can be divided into two categories:

- Query traffic – traffic between dispatch nodes N_α and search nodes N_β . This traffic is present all the time when searching is performed. The query traffic is characterized by low to medium data volumes and high frequency.
- Data traffic – traffic between indexing nodes N_γ and search nodes N_β . The data traffic is characterized by high data volumes and low frequency (typically one batch per day).

A typical query will transfer a query string from the dispatch node N_α to the search nodes N_β . Then the search nodes N_β will reply with a sequence of documents matching the query. Optionally N_α shall also be able to query the search node for the URL strings for the document, but this is considered immaterial in the present context.

The architecture of the search engine according to the invention can, based on the above considerations, now easily be scaled in two dimensions, viz. the data volume and the performance dimension respectively,

Data volume scaling is achieved by adding more data set partitions d , in other words more groups or columns S of search nodes N_β are added. Also the number of indexing nodes N_γ and dispatch nodes N_α can be increased as necessary in order to handle more data set partitions d .

Performance scaling can be achieved in the search engine architecture by replicating data set partitions d with a corresponding increase in the number of search nodes N_β , by increasing the number as illustrated in fig. 6. When using replication of data set partitions to scale the system performance, each search node N_β is part of the search node group S . Thus, the search nodes

$N_{\beta 1}, \dots, N_{\beta s}$ are arranged into groups $S_{\beta 1}, \dots, S_{\beta v}$, where $v = \frac{s}{h_s}$, h_s denoting the

scaling factor. The group $S_{\beta j}$ contains the search nodes $N_{\beta j}, N_{\beta j+1}, N_{\beta j+2}$ and $N_{\beta j+3}$ as rendered in fig. 8.

Scaling the data volume may cause the number of search nodes N_β receiving queries broadcast from a dispatch node N_α growing quite large. The architecture solves this problem by using several levels λ of dispatch nodes

N_α – this is illustrated in fig. 7. which renders the arrangement of the dispatch nodes N_α as nodes in a portion of a binary data distribution tree. A binary data distribution tree easily allows for a linear scalability. Similar binary data distribution trees of the kind that already has been disclosed in the applicant's above-mentioned international application PCT/NO99/00344, which discloses the configuration of an actual implementation of the pattern matching chip PMC. The number of dispatch nodes N_α in a regular binary tree is, of course $2^{\lambda-1}$ on each level λ , $\lambda = 1, 2, 3, \dots$. A dispatch root node is on the first level, and up to and including a given level λ there is a total of $2^\lambda - 1$ dispatch nodes in the tree. In case the dispatch nodes N_α are used also as acquisition nodes, i.e. for gathering the answers returned from the search nodes, the results of a search is merged in the dispatch nodes, the root dispatch node outputting the final answer to the query. However, there is nothing against that the search engine according to the invention is set up with a separate data gathering tree connected to the search nodes and comprising acquisition nodes N_δ gathering and outputting the final result of a query on the acquisition root node of the data gathering tree, i.e. the data acquisition node tree. The acquisition node tree could then be a mirror image of the dispatch node tree.

The schematic layout of a scalable search engine architecture according to the invention is shown in fig. 8, with the principle for two-dimensional scaling illustrated. It will be seen that the dispatch nodes N_α constitute a front end of the search engine and route the queries to the search node N_β and receives the search results back from the search nodes, wherein the actual search of the indexed data is performed. In case dedicated acquisition nodes N_δ are used, as shown in fig. 9, which otherwise is similar to fig. 8, the search results would, of course, be returned thereto. With the dispatch nodes N_α arranged in a tree configuration as shown in fig. 9, the acquisition node network as a back end of the search engine would form a mirror image of the dispatch node network. The indexing (spidering) nodes N_γ also constitute a back end of the search engine and collect data from e.g. the Internet and index the data to generate a searchable catalogue. By adding search nodes N_β or rather search nodes groups S horizontally the search engine scales linearly in the data volume, each additional search node or search node group containing different data. Typical capacity parameters for a search engine can as a non-limiting instance be given as follows. One search node N_β can

typically handle 8 000 000 page views per day in a 5 000 000 documents catalogue. For a scalable search engine each search node N_p typically could hold 5 000 000 unique indexed documents, implying that 40 search nodes in one row are enough to maintain a 200 000 000 documents catalogue. Scaling the performance, i.e. increasing the traffic capacity demands more rows of search nodes N_p with the same data to be added, such that the search nodes in a single column or group S contain identical data. A group or column S of 10 search nodes N_p hence will be able to handle 80 000 000 page views per day, with 40 columns handling a total of 3 200 000 000 page views per day.

An additional important benefit of a search engine according to the invention with the architecture scalable as herein disclosed, is that the query response time is essentially independent of the catalogue size, as each query is executed in parallel on all search nodes N_p and that the architecture is inherently fault tolerant, such that faults in the individual nodes will not result in a system breakdown, only temporary reduce the performance until the fault is corrected.

Moreover, the in principle unlimited linear scalability of the data volume and the traffic volume which can be provided in a search engine according to the invention, contrasts sharply with prior art search engines, wherein the search cost typically increases exponentially with the data or traffic volume increase, and wherein the maximum capacity of the prior art search engines typically will be reached at low to moderate volumes. With the search engine according to the invention the cost will scale linearly with the increase in capacity at most, depending actually on whether the capacity increase is provided by adding SW search nodes only or also SW/HW search nodes. Finally the search engine according to the invention offers the advantage that each node in practice can be implemented with standard low cost commercially available PCs, but alternatively also with more expensive UNIX-based servers such as for instance the Sun or Alpha computers as currently available.

PATENT CLAIMS

1. A search engine with two-dimensional linearly scalable parallel architecture for searching a collection of text documents D , wherein the documents can be divided into a number of partitions d_1, d_2, \dots, d_n , wherein the collection of documents D is preprocessed in a text filtration system such that a preprocessed document collection D_p is obtained and corresponding preprocessed partitions $d_{p1}, d_{p2}, \dots, d_{pn}$, wherein an index I can be generated from the document collection D such that for each previous preprocessed partition $d_{p1}, d_{p2}, \dots, d_{pn}$ a corresponding index i_1, i_2, \dots, i_n is obtained, wherein searching a partition d of the document collection D takes place with a partition-dependent data set $d_{p,k}$ comprising both the preprocessed partition d_{pk} and the corresponding index i_k , with $1 \leq k \leq n$, and wherein the search engine comprises data processing units which form sets of nodes (N) connected in a network, characterized in that
 - the first set of nodes comprises a dispatch nodes (N_{a1}, \dots, N_{aa}), a second set of nodes comprises b search nodes ($N_{\beta 1}, \dots, N_{\beta b}$), a third set of nodes comprises g indexing nodes ($N_{\gamma 1}, \dots, N_{\gamma g}$), and an optional fourth set of nodes comprises e acquisition nodes ($N_{\delta 1}, \dots, N_{\delta e}$).
 - that the dispatch nodes (N_a) are connected in a multilevel configuration in the network,
 - that the search nodes (N_β) are grouped in columns (S) which are connected in parallel in the network between the dispatch nodes (N_a) and an indexing node (N_γ), that the dispatch nodes (N_a) are adapted to process search queries and search answers, the search queries being dispatched further to all search nodes (N_β), and in case the acquisition nodes (N_δ) are not present, the search answers being returned to the dispatch nodes (N_a) and therein being combined to a final search result,
 - that the search nodes (N_β) each are adapted to contain search software,
 - that at least some of the search nodes (N_β) additionally comprise at least one search processor module (M),
 - that the indexing nodes (N_γ) are adapted for generally generating indexes i for the search software and optionally for generating partition-dependent data sets $d_{p,k}$ to search nodes (N_β) which comprise a search processor module,
 - that in case acquisition nodes (N_δ) are present, these are connected in a multilevel configuration in the network similar to that of the dispatch nodes

(N_a), and adapted for gathering answers to search queries and outputting a final result thereof, thus relieving the dispatch nodes of this task, and that the two-dimensional linear scaling respectively takes place by scaling of the data volume through an increase in the number of partitions d and scaling of performance through replication of one or more partitions d .

2. A search engine according to claim 1, characterized in that the multilevel configuration of the dispatch nodes (N_a) and the optional acquisition nodes (N_δ) in the network are provided by hierarchical tree structures.

3. A search engine according to claim 2, characterized in that multilevel configuration of the optional acquisition nodes (N_δ) is a mirror image of the multilevel configuration of the dispatch nodes (N_a).

4. A search engine according to claim 2, characterized in that the hierarchical tree structures are binary tree structures.

5. A search engine according to claim 1, characterized in that each of the search nodes (N_β) comprises a search software module (SW).



6. A search engine according to claim 5, characterized in that at least some of the search nodes (N_β) comprises at least one dedicated search processor module (M), each dedicated search processor module (M) being realized with one or more dedicated search processor chips (PMC) which each is adapted for parallel handling of a number of q search queries.

7. A search engine according to claim 6, characterized in that the dedicated search processor chips (PMC) are provided in the search processor modules (M) in y processor groups (G), each with z search processor chips (PMC) and being connected with and adapted to receive data from a memory (RAM) assigned to the processor group (G).

8. A search engine according to claim 1, characterized in that the increase in the number of partitions d in the scaling of the data volume is implemented by an increase in the number of search node groups or columns (S).

9. A search engine according to claim 8,
characterized in that an increase in the number of partitions d in the scaling
of the data volume is accompanied by a corresponding increase in the number
of dispatch nodes (N_α) and, in case, also in the number of acquisition nodes
5 (N_δ), and optionally also by an increase in the number of index nodes (N_γ).
10. A search engine according to claim 1,
characterized in that the replication of one or more partitions d in the scaling
of the performance is implemented by an increase in the number of search
nodes (N_β) in each group or column (S).
- 10 11. A search engine according to claim 1,
characterized in that the separate node sets (N_α , N_β , N_γ , N_δ) each is
implemented over one or more workstations connected in a data
communications network.

1/6



2/6




Fig. 3






Fig. 6

3/6



4/6




Fig.7

5/6




Fig.8

6/6




Fig.9

09/743268

Attorney Docket No.

3696-0101P

BIRCH, STEWART, KOLASCH & BIRCH, LLP

P.O. Box 747 • Falls Church, Virginia 22040-0747
Telephone: (703) 205-8000 • Facsimile: (703) 205-8050

PLEASE NOTE:
YOU MUST
COMPLETE THE
FOLLOWING

COMBINED DECLARATION AND POWER OF ATTORNEY FOR PATENT AND DESIGN APPLICATIONS

As a below named inventor, I hereby declare that: my residence, post office address and citizenship are as stated next to my name; that I verily believe that I am the original, first and sole inventor (if only one inventor is named below) or an original, first and joint inventor (if plural inventors are named below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

Insert Title: A SEARCH ENGINE WITH TWO-DIMENSIONAL LINEARLY SCALABLE PARALLEL ARCHITECTURE

Fill in Appropriate Information - For Use Without Specification Attached: the specification of which is attached hereto. If not attached hereto, the specification was filed on _____ as United States Application Number _____ and amended on _____ (if applicable) and/or the specification was filed on May 10, 2000 as PCT International Application Number PCT/NO00/00155; and was amended under PCT Article 19 on _____ (if applicable)

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to patentability as defined in Title 37, Code of Federal Regulations, §1.56.

I do not know and do not believe the same was ever known or used in the United States of America before my or our invention thereof, or patented or described in any printed publication in any country before my or our invention thereof or more than one year prior to this application, that the same was not in public use or on sale in the United States of America more than one year prior to this application, that the invention has not been patented or made the subject of an inventor's certificate issued before the date of this application in any country foreign to the United States of America on an application filed by me or my legal representative or assigns more than twelve months (six months for designs) prior to this application, and that no application for patent or inventor's certificate on this invention has been filed in any country foreign to the United States of America prior to this application by me or my legal representatives or assigns, except as follows.

I hereby claim foreign priority benefits under Title 35, United States Code, §119(a)-(d) of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Insert Priority Information: (if appropriate)

Prior Foreign Application(s)			Priority Claimed	
<u>19992269</u> (Number)	<u>Norway</u> (Country)	<u>May 10, 1999</u> (Month/Day/Year Filed)	<input checked="" type="checkbox"/> Yes	<input type="checkbox"/> No
_____ (Number)	_____ (Country)	_____ (Month/Day/Year Filed)	<input type="checkbox"/> Yes	<input type="checkbox"/> No
_____ (Number)	_____ (Country)	_____ (Month/Day/Year Filed)	<input type="checkbox"/> Yes	<input type="checkbox"/> No
_____ (Number)	_____ (Country)	_____ (Month/Day/Year Filed)	<input type="checkbox"/> Yes	<input type="checkbox"/> No

I hereby claim the benefit under Title 35, United States Code, §119(e) of any United States provisional applications(s) listed below.

Insert Provisional Application(s): (if any)

_____ (Application Number)	_____ (Filing Date)
_____ (Application Number)	_____ (Filing Date)

All Foreign Applications, if any, for any Patent or Inventor's Certificate Filed More than 12 Months (6 Months for Designs) Prior to the Filing Date of This Application:

Country	Application Number	Date of Filing (Month/Day/Year)
_____	_____	_____
_____	_____	_____

Insert Requested Information: (if appropriate)

I hereby claim the benefit under Title 35, United States Code, §120 of any United States and/or PCT application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States and/or PCT application in the manner provided by the first paragraph of Title 35, United States Code, §112, I acknowledge the duty to disclose information which is material to the patentability as defined in Title 37, Code of Federal Regulations, §1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application.

Insert Prior U.S. Application(s): (if any)

(Application Number)	(Filing Date)	(Status - patented, pending, abandoned)
_____	_____	_____

I hereby appoint the following attorneys to prosecute this application and/or an international application based on this application and to transact all business in the Patent and Trademark Office connected therewith and in connection with the resulting patent based on instructions received from the entity who first sent the application papers to the attorneys identified below, unless the inventor(s) or assignee provides said attorneys with a written notice to the contrary:

Raymond C. Stewart
Joseph A. Kolasch
Bernard L. Sweeney
Charles Gorenstein
Leonard R. Svensson
Andrew D. Meikle
Joe McKinney Muncy
John W. Bailey
Gary D. Yacura

(Reg. No. 21,066)
(Reg. No. 22,463)
(Reg. No. 24,448)
(Reg. No. 29,271)
(Reg. No. 30,330)
(Reg. No. 32,868)
(Reg. No. 32,334)
(Reg. No. 32,881)
(Reg. No. 35,416)

Terrell C. Birch
James M. Slattery
Michael K. Mutter
Gerald M. Murphy, Jr.
Terry L. Clark
Marc S. Weiner
Donald J. Daley
John A. Castellano

(Reg. No. 19,382)
(Reg. No. 28,380)
(Reg. No. 29,680)
(Reg. No. 28,977)
(Reg. No. 32,644)
(Reg. No. 32,181)
(Reg. No. 34,313)
(Reg. No. 35,094)

Send Correspondence to:

BIRCH, STEWART, KOLASCH & BIRCH, LLP

or **Customer No. 2292**

P.O. Box 747 • Falls Church, Virginia 22040-0747

Telephone: (703) 205-8000 • Facsimile: (703) 205-8050

PLEASE NOTE:
YOU MUST
COMPLETE
THE
FOLLOWING:

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of First
or Sole Inventor:
Insert Name of
Inventor
Insert Date This
Document is Signed

Insert Residence
Insert Citizenship

Insert Post Office
Address

Full Name of Second
Inventor, if any:
see above

Full Name of Third
Inventor, if any:
see above

Full Name of Fourth
Inventor, if any:
see above

Full Name of Fifth
Inventor, if any:
see above

GIVEN NAME/FAMILY NAME	INVENTOR'S SIGNATURE	DATE*
Børge SVINGEN	<i>Børge Svingen</i>	29.1.2001
Residence (City, State & Country)	CITIZENSHIP	
Hommelvik Norway	Norwegian	
POST OFFICE ADDRESS (Complete Street Address including City, State & Country)		
Borkdalen 6, N-7550 HOMMELVIK, NORWAY		
GIVEN NAME/FAMILY NAME	INVENTOR'S SIGNATURE	DATE*
Knut Magne RISVIK	<i>Knut Magne Risvik</i>	19.12.2000
Residence (City, State & Country)	CITIZENSHIP	
Trondheim, Norway	Norwegian	
POST OFFICE ADDRESS (Complete Street Address including City, State & Country)		
Sigrid Undsets vei 27, N-7023 TRONDHEIM, NORWAY		
GIVEN NAME/FAMILY NAME	INVENTOR'S SIGNATURE	DATE*
Arne HALAAS	<i>Arne Halaas</i>	20.12.2000
Residence (City, State & Country)	CITIZENSHIP	
Trondheim, Norway	Norwegian	
POST OFFICE ADDRESS (Complete Street Address including City, State & Country)		
Overlege Bratts vei 60, N-7025 TRONDHEIM, NORWAY		
GIVEN NAME/FAMILY NAME	INVENTOR'S SIGNATURE	DATE*
Tor EGGE	<i>Tor Egge</i>	8.1.2001
Residence (City, State & Country)	CITIZENSHIP	
Trondheim, Norway	Norwegian	
POST OFFICE ADDRESS (Complete Street Address including City, State & Country)		
Hjalmar Haalkes vei 1, N-7033 TRONDHEIM		
GIVEN NAME/FAMILY NAME	INVENTOR'S SIGNATURE	DATE*
Residence (City, State & Country)	CITIZENSHIP	
POST OFFICE ADDRESS (Complete Street Address including City, State & Country)		